arXiv:1401.4566v2 [cs.LG] 8 Feb 2014

Excess Risk Bounds for Exponentially Concave Losses

Mehrdad Mahdavi
Michigan State University
mahdavim@cse.msu.edu

Rong Jin
Michigan State University
rongjin@cse.msu.edu

Abstract

The overarching goal of this paper is to derive excess risk bounds for learning from exp-concave loss functions in passive and sequential learning settings. Exp-concave loss functions encompass several fundamental problems in machine learning such as squared loss in linear regression, logistic loss in classification, and negative logarithm loss in portfolio management. In batch setting, we obtain sharp bounds on the performance of empirical risk minimization performed in a linear hypothesis space and with respect to the exp-concave loss functions. We also extend the results to the online setting where the learner receives the training examples in a sequential manner. We propose an online learning algorithm that is a properly modified version of online Newton method to obtain sharp risk bounds. Under an additional mild assumption on the loss function, we show that in both settings we are able to achieve an excess risk bound of $O(d \log n/n)$ that holds with a *high probability*.

1 Introduction

We investigate the excess risk bounds for learning a linear classifier using a exponentially concave (abbr. as exp-concave) loss function (see e.g., [1] and [2]). More specifically, let $\mathcal{S} = \{(\mathbf{x}_1, y_1), (\mathbf{x}_2, y_2), \cdots, (\mathbf{x}_n, y_n)\} \in \Xi^n$ be a set of i.i.d. training examples sampled from an unknown distribution P over instance space $\Xi = \mathcal{X} \times \mathcal{Y}$, where $\mathbf{x}_i \in \mathcal{X} \subseteq \mathbb{R}^d$ with $\|\mathbf{x}_i\| \leq 1$ and $\mathbf{y}_i \in \mathcal{Y} := \{-1, +1\}$ and $\mathbf{y}_i \in \mathcal{Y} := [-1, +1]$ in classification and regression problems, respectively. Let $\mathcal{W} = \{\mathbf{w} \in \mathbb{R}^d : \|\mathbf{w}\| \leq R\}$ be our domain of linear classifiers with bounded norm, where $R > 0$ determines the size of the domain. We aim at finding a learner $\mathbf{w} \in \mathcal{W}$ with the assist of training samples $\mathcal{S}$ that generalizes well on unseen instances.

Let $\ell(z) : \mathbb{R} \mapsto \mathbb{R}_+$ be the convex surrogate loss function used to measure the classification error. In this work, we are interested in learning problems where the loss function $\ell(z)$ is a one-dimensional exponentially concave function with constant $\alpha > 0$ (i.e., $\exp(-\alpha z)$ is concave for any $|z| \leq R$). Examples of such loss functions are the squared loss used in regression, logistic loss used in classification, and negative logarithm loss used in portfolio management [3, 1, 4, 5]. Similar to most analysis of generalization performance, we assume $\ell(z)$ to be Lipschitz continuous with constant G, i.e. $|\ell'(z)| \leq G$. Define $\mathcal{L}(\mathbf{w})$ as the expected loss function for an arbitrary classifier $\mathbf{w} \in \mathcal{W}$, i.e.

$$\mathcal{L}(\mathbf{w}) = \mathrm{E}_{(\mathbf{x},y)\sim P}\left[\ell(y\mathbf{w}^\top \mathbf{x})\right].$$

Let $\mathbf{w}_* \in \mathcal{W}$ be the optimal solution that minimizes $\mathcal{L}(\mathbf{w})$ over the domain $\mathcal{W}$, i.e. $\mathbf{w}_* = \arg\min_{\mathbf{w}\in\mathcal{W}} \mathcal{L}(\mathbf{w})$. We note that the exp-concavity of individual loss functions $\ell(\cdot)$ also implies the exp-concavity of the expected function $\mathcal{L}(\mathbf{w})$ (a straightforward proof can be found in [3, Lemma 1]). Our goal is to efficiently learn a classifier $\widehat{\mathbf{w}}$ with the help of training set $\mathcal{S}$ with small excess risk defined by:

$$\mathcal{E}_P(\widehat{\mathbf{w}}) := \mathcal{L}(\widehat{\mathbf{w}}) - \min_{\mathbf{w}\in\mathcal{W}} \mathcal{L}(\mathbf{w}) = \mathcal{L}(\widehat{\mathbf{w}}) - \mathcal{L}(\mathbf{w}_*).$$

While the main focus of statistical learning theory was on understanding learnability and sample complexity by investigating the complexity of hypothesis class in terms of known combinatorial measures, recent advances in online learning and optimization theory opened a new trend in understanding the generalization ability of learning algorithms in terms of the characteristics of loss

functions being used in convex learning problems. In particular, a staggering number of results have focused on strong convexity of loss function (that is a stronger condition than exp-concavity) and obtained better generalization bounds which are referred to as *fast* rates [6, 7]. In terms of smoothness of loss function, a recent result [8] has shown that under smoothness assumption, it is possible to obtain *optimistic* rates (in the sense that smooth losses yield better generalization bounds when the problem is easier and the expected loss of optimal classifier is small), which are more appealing than Lipschitz continuous cases. This work extends the results to exp-concave loss functions and investigates how to obtain sharper excess risk bounds for learning from such functions. We note that although the online Newton method [1] yields $O(d \log n)$ regret bound, it is only able to achieve an $O(d \log n/n)$ bound for excess risk in expectation. In contrast, the excess risk bounds analyzed in this work are all in high probability sense.

We consider two settings to learn a classifier from the provided training set $\mathcal{S}$. In *statistical setting* (also called batch learning) [9], we assume that the learner has access to all training examples in advance, and in *online setting* the examples are assumed to become available to the learner one at a time. We show that with an additional assumption regarding the exponential concave loss function, we will be able to achieve an excess risk bound of $O(d \log n/n)$, which is significantly faster than $O(1/\sqrt{n})$ rate for general convex Lipschitz loss functions. The proof of batch setting utilizes the notion of local Radamacher complexities and involves novel ingredients tailored to exp-concave functions in order to obtain sharp convergence rates. In online setting, the results follows from Bernstein inequality for martingales and peeling process. We note that fast rates are possible and well known in sequential prediction via the notion of mixable losses [10], and in batch setting under Tsybakov's margin condition with $\kappa = 1$ [11], where the relation between these two settings has been recently investigated via the notion of stochastic mixability [12]. However, our analysis and conditions are different and only focuses on the exp-concavity property of the loss to derive an $O(\log n/n)$ risk bound.

2 The Algorithms

We study two algorithms for learning with exp-concave loss functions. The first algorithm that is devised for batch setting is simply based on empirical risk minimization. More specifically, it learns a classifier from the space of linear classifiers $\mathcal{W}$ by solving the following optimization problem

$$\min_{\mathbf{w} \in \mathcal{W}} \frac{1}{n} \sum_{i=1}^{n} \ell(y_i \mathbf{w}^\top \mathbf{x}_i). \tag{1}$$

The optimal solution to (1) is denoted by $\widehat{\mathbf{w}}_*$. Here, we are not concerned with the optimization procedure to find $\widehat{\mathbf{w}}_*$ and only investigate the access risk of obtained classifier $\widehat{\mathbf{w}}_*$ with respect to the optimal classifier $\mathbf{w}_*$.

Our second algorithm is a modified online Newton method [1]. Algorithm 1 gives the detailed steps. The key difference between Algorithm 1 and the online Newton algorithm [1] is that at each iteration, it estimates a smoothed version of the covariance matrix $\mathbf{Z}$ using the training examples received in the past. In contrast, the online Newton method takes into account the gradient $\ell'(y_i \mathbf{w}_i^\top \mathbf{x}_i)$ when updating $\mathbf{Z}_i$. It is this difference that allows us to derive an $O(d \log n/n)$ excess risk bound for the learned classifier. The classifier learned from the online algorithm $\widehat{\mathbf{w}}_*$ is simply the average of solutions obtained over all iterations. We also note that the idea of using an estimated covariance matrix for online learning and optimization has been examined by several studies [13, 14, 15, 16]. It is also closely related to the technique of time varying potential discussed in [17, 2] for regression. Unlike these studies that are mostly focused on obtaining regret bound, we aim to study the excess risk bound for the learned classifier.

3 Main Results

We first state the result of batch learning problem in (1), and then the result of online learning algorithm that is detailed in Algorithm 1. In order to achieve an excess risk bound better than

Algorithm 1 `Stochastic Exp-concave Optimization`

Input: step size $\eta_1 > 0$ and smoothing parameter $a > 0$
Initialization: $\mathbf{w}_1 = \mathbf{0}$ and $\mathbf{M}_0 = a\mathbf{I}$
for $i = 1, \ldots, n$ **do**
  Receive training example $(\mathbf{x}_i, y_i)$
  Compute $\mathbf{M}_i = \mathbf{M}_{i-1} + \mathbf{x}_i\mathbf{x}_i^\top$ and the covariance matrix $\mathbf{Z}_i = \mathbf{M}_i/i$
  Compute the gradient $\mathbf{v}_i = \ell'(y_i\mathbf{w}_i^\top\mathbf{x}_i)\mathbf{x}_i$ and step size $\eta_i = \eta_1/i$
  Update the solution $\mathbf{w}_i$ by solving the following optimization problem

$$\mathbf{w}_{i+1} = \arg\min_{\mathbf{w}\in\mathcal{W}} \eta_i\langle\mathbf{w}, \mathbf{v}_i\rangle + \frac{1}{2}\|\mathbf{w} - \mathbf{w}_i\|^2_{\mathbf{Z}_i} \tag{2}$$

  where $\|\mathbf{w}\|_{\mathbf{Z}_i} = \mathbf{w}^\top\mathbf{Z}_i\mathbf{w}$.
end for
return $\widehat{\mathbf{w}}_* = \frac{1}{n}\sum_{i=1}^n \mathbf{w}_i$.

$O(1/\sqrt{n})$, we introduce following key assumption for the analysis of the empirical error minimization problem in (1)

Assumption I there exists a constant $\theta > 0$ s. t. $\mathrm{E}\left[[\ell'(y\mathbf{w}_*^\top\mathbf{x})]^2\mathbf{x}\mathbf{x}^\top\right] \succeq \theta\mathrm{E}\left[\mathbf{x}\mathbf{x}^\top\right]$.

For the online learning method in Algorithm 1, we strengthen Assumption (I) as:

Assumption II there exists a constant $\theta > 0$ s. t. $\mathrm{E}\left[[\ell'(y\mathbf{w}^\top\mathbf{x})]^2\mathbf{x}\mathbf{x}^\top\right] \succeq \theta\mathrm{E}[\mathbf{x}\mathbf{x}^\top], \forall\mathbf{w} \in \mathcal{W}$.

Note that unlike Assumption (I) that only requires the property to hold with respect to the optimal solution $\mathbf{w}_*$, Assumption (II) requires the property to hold for any $\mathbf{w} \in \mathcal{W}$, making a stronger assumption than Assumption (I). We also note that Assumption (II) is closely related to strong convexity assumption. In particular, it is easy to verify that when $\mathrm{E}[\mathbf{x}\mathbf{x}^\top]$ is strictly positive definite, the expected loss $\mathcal{L}(\mathbf{w})$ will be strongly convex in $\mathbf{w}$ by using the property of exponential concave function.

The following lemma shows a general scenario when both Assumptions (I) and (II) hold.

Lemma 1. *Suppose (i)* $\Pr(y = 1|\mathbf{x}) \geq q$ *and* $\Pr(y = -1|\mathbf{x}) \geq q$ *for any* $\mathbf{x} \in \mathcal{X}$*, where* $q > 0$*, and (ii)* $\ell'(0) > 0$*. Then Assumption (I) and (II) hold with* $\theta \geq q[\ell'(0)]^2$

Proof. We first bound $\mathrm{E}\left[[\ell'(y\mathbf{w}^\top\mathbf{x})]^2|\mathbf{x}\right]$ for any given $\mathbf{x} \in \mathcal{X}$ by

$$\mathrm{E}\left[[\ell'(y\mathbf{w}^\top\mathbf{x})]^2|\mathbf{x}\right] \geq q\left([\ell'(\mathbf{w}^\top\mathbf{x})]^2 + [\ell'(-\mathbf{w}^\top\mathbf{x})]^2\right)$$

Since $\ell'(z)$ is monotonically increasing function, we have

$$|\ell'(0)| \leq \max(|\ell'(\mathbf{w}^\top\mathbf{x})|, |\ell'(-\mathbf{w}^\top\mathbf{x})|)$$

and therefore

$$\mathrm{E}\left[[\ell'(y\mathbf{w}^\top\mathbf{x})]^2|\mathbf{x}\right] \geq q[\ell'(0)]^2$$

implying that $\mathrm{E}\left[[\ell'(y\mathbf{w}^\top\mathbf{x})]^2\mathbf{x}\mathbf{x}^\top\right] \geq q[\ell'(0)]^2\mathrm{E}[\mathbf{x}\mathbf{x}^\top]$ as desired. □

We note that $\ell'(0) > 0$ is the necessary and sufficient condition that the convex surrogate loss function for 0-1 loss function to be classification-calibrated [18], and therefore is almost unavoidable if our final goal is to minimize the binary classification error.

The excess risk bound for the batch learning algorithm is given in the following theorem.

Theorem 1. *Suppose Assumption (I) holds. Let* $\widehat{\mathbf{w}}_*$ *be the solution to the convex optimization problem in (1). Define*

$$\gamma = \max\left(1, \frac{G^2}{\theta}, \frac{G}{\alpha\theta R}\right), \quad \rho_0 = \max\left(32\gamma, \sqrt{\gamma\left(28 + \frac{3}{GR}\right)}\right).$$

Then with a probability $1 - 2me^{-t}$*, where* $m = \lceil \log_2 n \rceil$*, we have*

$$\mathcal{L}(\widehat{\mathbf{w}}_*) - \mathcal{L}(\mathbf{w}_*) \leq \left(GR\left[32\rho_0 + 28\right] + 3\right) \frac{t + 2 + d\log n}{n} = \tilde{O}\left(\frac{d\log n}{n}\right).$$

The following theorem provides the excess risk bound for Algorithm 1 where training examples ae received in an online fashion and the final solution is reported as the average of all the intermediate solution.

Theorem 2. *Suppose Assumption (II) holds. Let* $\widehat{\mathbf{w}}_*$ *be the average solution returned by Algorithm 1, with* $\eta_1 = \max(1, 3/[\theta\beta])$ *and* $a = \eta_1^2 G^2 d/[4R^2]$*. With a probability* $1 - 2me^{-t}$*, where* $m = \lceil \log_2 n \rceil$*, we have*

$$\mathcal{L}(\widehat{\mathbf{w}}_*) - \mathcal{L}(\mathbf{w}_*) \leq \rho_0 GR\frac{d}{n}\log\left(1 + \frac{4n}{\gamma^2 d^2}\right) + 3GR(2\gamma + 1)t = \tilde{O}\left(\frac{d\log n}{n}\right).$$

Remark 1. *As indicated in Theorems 1 and 2, the excess risk for both batch learning and online learning is reduced at the rate of* $O(d\log n/n)$*, which is consistent with the regret bound for online optimizing the exponentially concave loss functions [1]. We note that the linear dependence on* d *is in general unavoidable. This is because when* $\mathrm{E}[\mathbf{x}\mathbf{x}^\top]$ *is strictly positive definite, the function* $\mathcal{L}(\mathbf{w})$ *will be strongly convex with modulus proportion to* $\lambda_{\min}\left(\mathrm{E}[\mathbf{x}\mathbf{x}^\top]\right)$*. Since* $\lambda_{\min}\left(\mathrm{E}[\mathbf{x}\mathbf{x}^\top]\right) \propto 1/d$*, we would expect a linear dependence on* d *based on the minimax convergence rate of stochastic optimization for strongly convex function. Finally, we note that for strongly convex loss functions, it is known that an* $O(1/n)$ *excess risk bound can be achieved without the* $\log n$ *factor. It is however unclear if the* $\log n$ *factor can be removed from the excess risk bounds for exponential concave functions, a question to be investigated in the future.*

Comparing the result of online learning with that of batch learning, we observe that, although both achieve similar excess risk bounds, batch learning algorithm is advantageous in two aspects. First, the batch learning algorithm has to make a weaker assumption about the data (i.e. Assumption (I) vs. Assumption (II)). Second, the batch learning algorithm does not have to know the parameter α *and* θ *in advance, which is important for online learning method to determine the step size* η_1*.*

4 Analysis

We now turn to the proofs of our main results. The main steps in each proof are provided in the main text, with some of the more technical results deferred to the appendix.

4.1 Proof of Theorem 1

Our analysis for batch setting is based on the Talagrand's inequality and in particular its variant (Klein-Rio bound) with improved constants derived in [19] (see also [20, Chapter 2]). To do so, we define

$$\|P_n - P\|_{\mathcal{W}} = \sup_{\mathbf{w}\in\mathcal{W}} \left| \frac{1}{n}\sum_{i=1}^{n} \left[\ell(y_i\mathbf{w}^\top\mathbf{x}_i) - \ell(y_i\mathbf{w}_*^\top\mathbf{x}_i)\right] - \mathrm{E}_{(\mathbf{x},y)}\left[\ell(y\mathbf{w}^\top\mathbf{x}) - \ell(y\mathbf{w}_*^\top\mathbf{x})\right] \right|$$

and

$$U(\mathcal{W}) = \max_{\mathbf{w}\in\mathcal{W}, \|\mathbf{x}\|\leq 1} \ell(y\mathbf{w}^\top\mathbf{x}) - \ell(y\mathbf{w}_*^\top\mathbf{x}),\ \sigma_P(\mathcal{W}) = \sup_{\mathbf{w}\in\mathcal{W}} \mathrm{E}_{(\mathbf{x},y)}\left[(\ell(y\mathbf{w}^\top\mathbf{x}) - \ell(y\mathbf{w}_*^\top\mathbf{x}))^2\right].$$

The analysis is rooted in the following concentration inequality:

Theorem 3. *We have*

$$\Pr\left\{\|P_n - P\|_{\mathcal{W}} \geq 2\mathrm{E}\|P_n - P\|_{\mathcal{W}} + \sigma_P(\mathcal{W})\sqrt{\frac{2t}{n}} + \frac{(U(\mathcal{W}) + 3)t}{3n}\right\} \leq e^{-t}.$$

The following property of exponential concave loss function from [1] will be used throughout the paper.

Theorem 4. *If a function $f : \mathcal{W} \mapsto \mathbb{R}$ is such that $exp(-\alpha f(\mathbf{w}))$ is concave, and has gradient bounded by $\|\nabla f\| \leq G$, then there exists $\beta \leq \frac{1}{2}\min(\alpha, 1/[4GR])$ such that the following holds*

$$f(\mathbf{w}) \geq f(\mathbf{w}') + (\mathbf{w} - \mathbf{w}')^\top \nabla f(\mathbf{w}') + \frac{\beta}{2}\left[\nabla f(\mathbf{w}')^\top (\mathbf{w} - \mathbf{w}')\right]^2, \forall \mathbf{w}, \mathbf{w}' \in \mathcal{W}.$$

The key quantity for our analysis is the following random variable:

$$\rho(\mathbf{w}) = \frac{1}{2R}\sqrt{\mathrm{E}[|\mathbf{x}^\top(\mathbf{w} - \mathbf{w}_*)|]^2}.$$

Evidently, $\rho(\mathbf{w}) \leq 1, \forall \mathbf{w} \in \mathcal{W}$. The following lemma deals with the concentration of $\rho(\mathbf{w})$.

Lemma 2. *Define $\Delta = \{\mathbf{w} \in \mathcal{W} : \rho(\mathbf{w}) \leq \rho\}$. Then, with a probability $1 - e^{-t}$, we have,*

$$\frac{1}{n}\sup_{\mathbf{w}\in\Delta}\sum_{i=1}^n \left[\mathbf{x}_i^\top(\mathbf{w} - \mathbf{w}_*)\right]^2 \leq 10R^2\left(\rho^2 + \frac{t + 1 + d\log n}{n}\right)$$

Proof. Fix a $\mathbf{w} \in \Delta := \{\mathbf{w} \in \mathcal{W} : \rho(\mathbf{w}) \leq \rho\}$. Using the standard Bernstein's inequality [21], we have, with a probability $1 - e^{-t}$,

$$\left|\frac{1}{n}\sum_{i=1}^n [\mathbf{x}_i^\top(\mathbf{w} - \mathbf{w}_*)]^2 - \mathrm{E}[((\mathbf{w} - \mathbf{w}_*)^\top \mathbf{x})^2]\right| \leq \frac{16R^2 t}{3n} + 2R\sqrt{\frac{2\mathrm{E}[((\mathbf{w} - \mathbf{w}_*)^\top \mathbf{x})^2]t}{n}}$$

By definition of the domain Δ, i.e., $\rho(\mathbf{w}) = \frac{1}{2R}\sqrt{\mathrm{E}[|\mathbf{x}^\top(\mathbf{w} - \mathbf{w}_*)|]^2} \leq \rho$ and above concentration result we obtain:

$$\frac{1}{n}\sum_{i=1}^n [\mathbf{x}_i^\top(\mathbf{w} - \mathbf{w}_*)]^2 \leq 4R^2\left(\rho^2 + \frac{4t}{3n} + \rho\sqrt{\frac{2t}{n}}\right) \leq 10R^2\left(\rho^2 + \frac{t}{n}\right)$$

Next, we consider a discrete version of the space Δ. Let $\mathcal{N}(\Delta, \epsilon)$ be the proper ϵ-net of Δ. Since $\Delta \subseteq \mathcal{W}$, we have

$$|\mathcal{N}(\Delta, \epsilon)| \leq |\mathcal{N}(\mathcal{W}, \epsilon)| \leq \left(\frac{3R}{\epsilon}\right)^d$$

Using the union bound, we have, with a probability $1 - e^{-t}$, for any $\mathbf{w} \in \mathcal{N}(\Delta, 3R/\sqrt{n})$,

$$\frac{1}{n}\sum_{i=1}^n [\mathbf{x}_i^\top(\mathbf{w} - \mathbf{w}_*)]^2 \leq 10R^2\left(\rho^2 + \frac{t + d\log n}{n}\right)$$

Since for any $\mathbf{w} \in \Delta$, there exists $\mathbf{w}' \in \mathcal{N}(\Delta, 3R/\sqrt{n})$, such that $\|\mathbf{w} - \mathbf{w}'\|_2 \leq 3R/\sqrt{n}$, we have, with a probability $1 - e^{-t}$, for any $\mathbf{w} \in \Delta$,

$$\frac{1}{n}\sum_{i=1}^n [\mathbf{x}_i^\top(\mathbf{w} - \mathbf{w}_*)]^2 \leq 10R^2\left(\rho^2 + \frac{t + 1 + d\log n}{n}\right),$$

as desired.

□

Define $\widehat{\rho} = \rho(\widehat{\mathbf{w}}_*)$. The next theorem allows us to bound the excess risk using the random variable $\widehat{\rho}$.

Theorem 5. *With a probability $1 - 2me^{-t}$, where $m = \lceil \log n \rceil$, we have*

$$\mathcal{L}(\widehat{\mathbf{w}}_*) - \mathcal{L}(\mathbf{w}_*) \leq GR\left(26\left[\widehat{\rho}\sqrt{\frac{\widetilde{t}}{n}} + \frac{\widetilde{t}}{n}\right] + 6\widehat{\rho}\sqrt{\frac{\widetilde{t}}{n}} + \frac{2\widetilde{t}}{n}\right) + \frac{3\widetilde{t}}{n}$$

where $\widetilde{t} = t + 2 + d\log n$.

Taking this statement as given for the moment, we proceed with the proof of Theorem 1, returning later to establish the claim stated in Theorem 5. Our overall strategy of proving Theorem 1 is to first bound $\widehat{\rho}$ by using the property of exp-concave function and the result from Theorem 5, and then bound the excess risk. More specifically, using the result from Theorem 5, we have, with a probability at least $1 - 2me^{-t}$,

$$\mathcal{L}(\widehat{\mathbf{w}}_*) - \mathcal{L}(\mathbf{w}_*) \leq GR \left(26 \left[\widehat{\rho} \sqrt{\frac{\widetilde{t}}{n}} + \frac{\widetilde{t}}{n} \right] + 6\widehat{\rho} \sqrt{\frac{\widetilde{t}}{n}} + \frac{2\widetilde{t}}{n} \right) + \frac{3\widetilde{t}}{n} \tag{3}$$

Using the property of exp-concave loss functions stated in Theorem 4, we have

$$\begin{aligned} \mathcal{L}(\widehat{\mathbf{w}}_*) - \mathcal{L}(\mathbf{w}_*) &\geq (\widehat{\mathbf{w}}_* - \mathbf{w}_*)^\top \nabla \mathcal{L}(\mathbf{w}_*) + \frac{\beta}{2} \mathrm{E}\left[[\ell'(y\mathbf{w}_*^\top \mathbf{x})(\widehat{\mathbf{w}}_* - \mathbf{w}_*)^\top \mathbf{x}]^2 \right] \\ &\geq \frac{\beta}{2} \mathrm{E}\left[[\ell'(y\mathbf{w}_*^\top \mathbf{x})(\widehat{\mathbf{w}}_* - \mathbf{w}_*)^\top \mathbf{x}]^2 \right] \end{aligned}$$

where the second step follows from the fact that $\mathbf{w}_*$ minimizes $\mathcal{L}(\mathbf{w})$ over the domain $\mathcal{W}$ and as a result $(\widehat{\mathbf{w}}_* - \mathbf{w}_*)^\top \nabla \mathcal{L}(\mathbf{w}_*) \geq 0$. We then use Assumption (I) to get

$$\mathrm{E}\left[[\ell'(y\mathbf{w}_*^\top \mathbf{x})(\widehat{\mathbf{w}}_* - \mathbf{w}_*)^\top \mathbf{x}]^2 \right] \geq \theta \mathrm{E}\left[((\widehat{\mathbf{w}}_* - \mathbf{w}_*)^\top \mathbf{x})^2 \right] = 4\theta R^2 \widehat{\rho}^2$$

and therefore

$$\mathcal{L}(\widehat{\mathbf{w}}_*) - \mathcal{L}(\mathbf{w}_*) \geq 2\beta\theta R^2 \widehat{\rho}^2. \tag{4}$$

Combining the bounds in (3) and (4), we have, with a probability $1 - 2me^{-t}$,

$$\widehat{\rho}^2 \leq \frac{G}{2\beta\theta R} \left(32\widehat{\rho} \sqrt{\frac{\widetilde{t}}{n}} + 28\frac{\widetilde{t}}{n} \right) + \frac{3\widetilde{t}}{2\beta\theta R^2 n}$$

implying that

$$\widehat{\rho} \leq \max\left(\frac{32G}{\beta\theta R}, \sqrt{\frac{G}{\beta\theta R}\left(28 + \frac{3}{GR} \right)} \right) \sqrt{\frac{\widetilde{t}}{n}}.$$

We derive the final bound for $\widehat{\rho}$ by plugging the bound for β. The excess risk bound is completed by plugging the above bound for $\widehat{\rho}$.

We now turn to proving the result stated in Theorem 5.

of Theorem 5. Our analysis will be based on the technique of local Rademacher complexity [22, 23, 20]. The notion of local Rademacher complexity works by considering Rademacher averages of smaller subsets of the hypothesis set. It generally leads to sharper learning bounds which, under certain general conditions, guarantee a faster convergence rate. Define $\rho_0 = 1/n$. We divide the range $[\rho_0, 1]$ into $m = \lceil \log_2 n \rceil$ segments, with $[\rho_0, \rho_1]$, $[\rho_1, \rho_2]$, ..., $[\rho_{m-1}, \rho_m]$, where $\rho_k = \rho_0 2^k$. Let $\widehat{\rho} = \rho(\widehat{\mathbf{w}}_*)$. Note that $\widehat{\rho}$ is a random variable depending on the sampled training examples.

As the first step, we assume that $\widehat{\rho} \in [\rho_k, \rho_{k+1}]$ for some fixed k. Define domain Δ as

$$\Delta = \{\mathbf{w} \in \mathcal{W} : \rho(\mathbf{w}) \leq \rho_{k+1}\}$$

Using the Telegrand inequality, with a probability at least $1 - e^{-t}$, we have

$$\mathcal{L}(\widehat{\mathbf{w}}_*) - \mathcal{L}(\mathbf{w}_*) \leq 2\mathrm{E}\|P_n - P\|_\Delta + \sigma_P(\Delta)\sqrt{\frac{2t}{n}} + \frac{(U(\Delta) + 3)\, t}{n} \tag{5}$$

We now bound each item on the right hand side of (5). First, we bound $\mathrm{E}\|P_n - P\|_\Delta$ as

$$\begin{aligned} \mathrm{E}\|P_n - P\|_\Delta &= \frac{2}{n} \mathrm{E}\left[\sup_{\mathbf{w} \in \Delta} \sum_{i=1}^n \sigma_i \left(\ell(y_i \mathbf{w}^\top \mathbf{x}_i) - \ell(y_i \mathbf{w}_*^\top \mathbf{x}_i) \right) \right] \\ &\leq \frac{4G}{n} \mathrm{E}\left[\sup_{\mathbf{w} \in \Delta} \sum_{i=1}^n y_i \sigma_i \mathbf{x}_i^\top (\mathbf{w} - \mathbf{w}_*) \right] \end{aligned}$$

where $\sigma_1, \ldots, \sigma_n$ are Rademacher random variables and the second step utilizes the contraction property of Rademacher complexity.

To bound $\mathrm{E}\|P_n - P\|_\Delta$, we need to bound $\sup_{\mathbf{w}\in\Delta} \sum_{i=1}^n [\mathbf{x}_i^\top(\mathbf{w} - \mathbf{w}_*)]^2$. Using Lemma 2, we have, with a probability $1 - e^{-t}$,

$$\begin{aligned}\mathrm{E}\|P_n - P\|_\Delta &\leq \frac{4G}{\sqrt{n}}\sqrt{\sup_{\mathbf{w}\in\Delta}\sum_{i=1}^n [\mathbf{x}_i^\top(\mathbf{w} - \mathbf{w}_*)]^2} \leq \frac{13GR}{\sqrt{n}}\sqrt{\rho_{k+1}^2 + \frac{t + 1 + d\log n}{n}} \\ &\leq 13GR\left(\frac{\sqrt{t + 1 + d\log n}}{n} + \frac{\rho_{k+1}}{\sqrt{n}}\right)\end{aligned}$$

Next, we bound $\sigma_P(\Delta)$ and $U(\Delta)$, i.e.

$$\begin{aligned}&\sigma_P^2(\Delta) \\ &\leq \sup_{\mathbf{w}\in\Delta} \mathrm{E}\left[(\ell(y\mathbf{w}^\top\mathbf{x}) - \ell(y\mathbf{w}_*^\top\mathbf{x}))^2\right] \\ &\leq \sup_{\mathbf{w}\in\Delta} G^2\mathrm{E}[((\mathbf{w} - \mathbf{w}_*)^\top\mathbf{x})^2] = 4R^2G^2\rho_{k+1}^2\end{aligned}$$

and $U(\Delta) \leq 2GR$. By putting the above results together, under the assumption $\widehat{\rho} \in [\rho_k, \rho_{k+1}]$, we have, with a probability $1 - 2e^{-t}$,

$$\mathcal{L}(\widehat{\mathbf{w}}_*) - \mathcal{L}(\mathbf{w}_*) \leq 26GR\left[\frac{\rho_{k+1}}{\sqrt{n}} + \frac{\sqrt{t + 1 + d\log n}}{n}\right] + 2GR\rho_{k+1}\sqrt{\frac{2t}{n}} + \frac{(2GR + 3)t}{n} \tag{6}$$

Define $\widetilde{t} = t + 2 + d\log n$. Using the fact that $\rho_{k+1} \leq 2\widehat{\rho}$, we can rewrite the bound in (6) as

$$\mathcal{L}(\widehat{\mathbf{w}}_*) - \mathcal{L}(\mathbf{w}_*) \leq 26GR\left[\frac{\widehat{\rho}\widetilde{t}}{\sqrt{n}} + \frac{\sqrt{\widetilde{t}}}{n}\right] + 6GR\widehat{\rho}\sqrt{\frac{\widetilde{t}}{n}} + \frac{(2GR + 3)\widetilde{t}}{n}$$

By taking the union bound over all the segments, with probability $1 - 2me^{-t}$, for any $\widehat{\rho} \in [\rho_0, 1]$, we have

$$\mathcal{L}(\widehat{\mathbf{w}}_*) - \mathcal{L}(\mathbf{w}_*) \leq 26GR\left[\widehat{\rho}\sqrt{\frac{\widetilde{t}}{n}} + \frac{\sqrt{\widetilde{t}}}{n}\right] + 6GR\widehat{\rho}\sqrt{\frac{\widetilde{t}}{n}} + \frac{(2GR + 3)\widetilde{t}}{n} \tag{7}$$

Finally, when $\widehat{\rho} \leq \rho_0 = 1/n$, we obtain

$$\mathcal{L}(\widehat{\mathbf{w}}_*) - \mathcal{L}(\mathbf{w}_*) \leq 2R\rho_0 G \leq \frac{2RG}{n} \tag{8}$$

We complete the proof by combining the bounds in (7) and (8). □

4.2 Proof of Theorem 2

We now turn to proving the main result on the excess risk for online setting. Define the covariance matrix $\mathbf{H}$ as $\mathbf{H} = \mathrm{E}_\mathbf{x}[\mathbf{x}\mathbf{x}^\top]$. The following theorem bounds $\mathcal{L}(\mathbf{w}_i) - \mathcal{L}(\mathbf{w}_*)$ by exploiting the property of exponentially concave functions (i.e., Theorem 4) and Assumption (II). Define δ_i as

$$\delta_i = \nabla\mathcal{L}(\mathbf{w}_i) - \ell'(y_i\mathbf{x}_i^\top\mathbf{w}_i)\mathbf{x}_i. \tag{9}$$

Lemma 3. *Suppose Assumption (II) holds. We have*

$$\begin{aligned}&\mathcal{L}(\mathbf{w}_i) - \mathcal{L}(\mathbf{w}_*) + \frac{\theta\beta}{3}\|\mathbf{w}_i - \mathbf{w}_*\|_\mathbf{H}^2 \\ &\leq \frac{\|\mathbf{w}_i - \mathbf{w}_*\|_{\mathbf{M}_{i-1}}^2}{2\eta_1} - \frac{\|\mathbf{w}_{i+1} - \mathbf{w}_*\|_{\mathbf{M}_i}^2}{2\eta_1} + \frac{\eta_1 G^2}{2}\mathbf{x}_i^\top\mathbf{M}_i^{-1}\mathbf{x}_i + (\mathbf{w}_i - \mathbf{w}_*)^\top\delta_i \\ &\quad + \frac{\theta\beta}{6}\left(\left[(\mathbf{w}_i - \mathbf{w}_*)^\top\mathbf{x}_i\right]^2 - \|\mathbf{w}_i - \mathbf{w}_*\|_\mathbf{H}^2\right),\end{aligned} \tag{10}$$

where $\|\mathbf{w}\|_\mathbf{H}^2 = \langle\mathbf{w}, \mathbf{H}\mathbf{w}\rangle$.

Lemma 4. *We have*

$$\mathbf{x}_i^\top \mathbf{M}_i^{-1} \mathbf{x}_i \leq \ln\det(\mathbf{M}_i) - \ln\det(\mathbf{M}_{i-1})$$

By using Lemma 4 and adding the inequalities in (10) over all the iterations, we have

$$\begin{aligned}
&\sum_{i=1}^n \mathcal{L}(\mathbf{w}_i) - \mathcal{L}(\mathbf{w}_*) + \frac{\theta\beta}{3}\sum_{i=1}^n \|\mathbf{w}_i - \mathbf{w}_*\|_{\mathbf{H}}^2 \\
&\leq \frac{\|\mathbf{w}_1 - \mathbf{w}_*\|_{\mathbf{M}_0}^2}{2\eta_1} + \frac{\eta_1 G^2}{2}\underbrace{\left(\log\det(\mathbf{M}_n) - \log\det(\mathbf{M}_0)\right)}_{\equiv\Delta_1} + \underbrace{\sum_{i=1}^n (\mathbf{w}_i - \mathbf{w}_*)^\top \delta_i}_{\equiv\Delta_2} \\
&\quad + \frac{\theta\beta}{6}\underbrace{\sum_{i=1}^n \left[(\mathbf{w}_i - \mathbf{w}_*)^\top \mathbf{x}_i\right]^2 - \|\mathbf{w}_i - \mathbf{w}_*\|_{\mathbf{H}}^2}_{\equiv\Delta_3}. \qquad (11)
\end{aligned}$$

We will bound Δ_1, Δ_2, and Δ_3, separately. We start by bounding Δ_1 as indicated by the following lemma.

Lemma 5.

$$\Delta_1 \leq d\log\left(1 + \frac{n}{ad}\right)$$

To bound Δ_2, we define $A = \sum_{i=1}^n \|\mathbf{w}_i - \mathbf{w}_*\|_{\mathbf{H}}^2$. Using the Berstein inequality for martingale [21] and peeling process [20], we have the following lemmas for bounding Δ_2 and Δ_3.

Lemma 6. *We have*

$$\Pr\left(A \leq \frac{4R^2}{n}\right) + \Pr\left(\Delta_2 \leq \left[\frac{6G^2}{\theta\beta} + GR\right]t + \frac{\theta\beta}{6}A\right) \geq 1 - me^{-t}$$

where $m = \lceil 2\log_2 n \rceil$.

Lemma 7. *We have*

$$\Pr\left(A \leq \frac{4R^2}{n}\right) + \Pr\left(\Delta_3 \leq 8R^2 t + A\right) \geq 1 - me^{-t}$$

where $m = \lceil 2\log_2 n \rceil$.

First, we consider the case when $A \leq 4R^2/n$ and show the following bound.

Lemma 8. *Assume that the condition $A \leq 4R^2/n$ holds. We have*

$$\sum_{i=1}^n \mathcal{L}(\mathbf{w}_i) - \mathcal{L}(\mathbf{w}_*) + \frac{\theta\beta}{2}\sum_{i=1}^n \|\mathbf{w}_i - \mathbf{w}_*\|_{\mathbf{H}}^2 \leq 2RG. \qquad (12)$$

Second, we assume that the following two conditions hold

$$\Delta_2 \leq \left[\frac{6G^2}{\theta\beta} + GR\right]t + \frac{\theta\beta}{6}A,\ \Delta_3 \leq 8R^2 t + A$$

Combining the above conditions with Lemma 5 and using the inequality in (11), we have

$$\sum_{i=1}^n \mathcal{L}(\mathbf{w}_i) - \mathcal{L}(\mathbf{w}_*) \leq \frac{2aR^2}{\eta_1} + \frac{\eta_1 G^2}{2} d\log\left(1 + \frac{n}{ad}\right) + \left[\frac{6G^2}{\theta\beta} + GR + \frac{4\theta\beta}{3}R^2\right]t.$$

Using the fact that $\eta_1 \geq 3/[\theta\beta]$, we set $a = \eta_1^2 G^2 d/[4R^2]$, we have

$$\sum_{i=1}^n \mathcal{L}(\mathbf{w}_i) - \mathcal{L}(\mathbf{w}_*) \leq \frac{3G^2}{\theta\beta} d\log\left(1 + \frac{4R^2\theta\beta^2 n}{G^2 d^2}\right) + \left[\frac{6G^2}{\theta\beta} + GR + \frac{4\theta\beta}{3}R^2\right]t.$$

We complete the proof by combining the two cases.

5 Conclusions and Future Work

In this work, we addressed the generalization ability of learning from exp-concave loss functions in batch and online settings. For both cases we show that the excess risk bound can be bounded by $O(d \log n/n)$ when the learning is performed in a linear hypothesis space with dimension d and with the help of n training examples.

One open question to be addressed in the future is if $\log n$ factor can be removed from the excess risk bound for exponentially concave loss functions by a more careful analysis. Another open question that needs to be investigated in the future is to improve the dependence on d if we are after a sparse solution. According to the literature of sparse recovery [24] and optimization [25], we should be able to replace d with $s \log d$ in the excess risk bound if we restrict the optimal solution to a sparse one. In the future, we plan to explore the technique of sparse recovery in analyzing the generalization performance of exponential concave function to reduce the dependence on d.

References

[1] E. Hazan, A. Agarwal, and S. Kale, "Logarithmic regret algorithms for online convex optimization," *Mach. Learn.*, vol. 69, no. 2-3, pp. 169–192, 2007.

[2] N. Cesa-Bianchi and G. Lugosi, *Prediction, Learning, and Games*. Cambridge University Press, 2006.

[3] T. Koren, "Open problem: Fast stochastic exp-concave optimization," in *COLT*, 2013.

[4] H. B. McMahan and M. J. Streeter, "Open problem: Better bounds for online logistic regression," in *COLT*, pp. 44.1–44.3, 2012.

[5] A. Agarwal, E. Hazan, S. Kale, and R. E. Shapire, "Algorithms for portfolio management based on the newton method," in *Proceedings of the 23rd International Conference on Machine Learning (ICML 2006)*, pp. 9–16, 2006.

[6] S. M. Kakade and A. Tewari, "On the generalization ability of online strongly convex programming algorithms," in *Advances in Neural Information Processing Systems*, pp. 801–808, 2008.

[7] K. Sridharan, S. Shalev-Shwartz, and N. Srebro, "Fast rates for regularized objectives," in *Advances in Neural Information Processing Systems*, pp. 1545–1552, 2008.

[8] N. Srebro, K. Sridharan, and A. Tewari, "Smoothness, low noise and fast rates," in *Advances in Neural Information Processing Systems*, pp. 2199–2207, 2010.

[9] O. Bousquet, S. Boucheron, and G. Lugosi, "Introduction to statistical learning theory," in *Advanced Lectures on Machine Learning*, pp. 169–207, Springer, 2004.

[10] V. Vovk, "A game of prediction with expert advice," in *Proceedings of the eighth annual conference on Computational learning theory*, pp. 51–60, ACM, 1995.

[11] A. B. Tsybakov, "Optimal aggregation of classifiers in statistical learning," *The Annals of Statistics*, vol. 32, no. 1, pp. 135–166, 2004.

[12] T. Van Erven, P. D. Grünwald, M. D. Reid, R. C. Williamson, *et al.*, "Mixability in statistical learning," in *Advances in Neural Information Processing Systems 25 (NIPS 2012)*, 2012.

[13] K. Crammer and D. D. Lee, "Learning via gaussian herding," in *Advances in neural information processing systems*, pp. 451–459, 2010.

[14] K. Crammer, A. Kulesza, and M. Dredze, "Adaptive regularization of weight vectors," *Machine Learning*, pp. 1–33, 2009.

[15] J. Duchi, E. Hazan, and Y. Singer, "Adaptive subgradient methods for online learning and stochastic optimization," *The Journal of Machine Learning Research*, vol. 11, pp. 2121–2159, 2011.

[16] F. Orabona and K. Crammer, "New adaptive algorithms for online classification," in *Advances in Neural Information Processing Systems*, 2010.

[17] N. Cesa-Bianchi, A. Conconi, and C. Gentile, "A second-order perceptron algorithm," in *Computational Learning Theory*, pp. 121–137, Springer, 2002.

[18] P. L. Bartlett, M. I. Jordan, and J. D. McAuliffe, "Convexity, classification, and risk bounds," *Journal of the American Statistical Association*, vol. 101, no. 473, pp. 138–156, 2006.

[19] T. Klein and E. Rio, "Concentration around the mean for maxima of empirical processes," *The Annals of Probability*, vol. 33, no. 3, pp. 1060–1077, 2005.

[20] V. Koltchinskii, *Oracle Inequalities in Empirical Risk Minimization and Sparse Recovery Problems*. Springer, 2011.

[21] S. Boucheron, G. Lugosi, and O. Bousquet, "Concentration inequalities," in *Advanced Lectures on Machine Learning*, pp. 208–240, Springer, 2004.

[22] P. L. Bartlett, O. Bousquet, and S. Mendelson, "Local rademacher complexities," *The Annals of Statistics*, vol. 33, no. 4, pp. 1497–1537, 2005.

[23] V. Koltchinskii, "Local rademacher complexities and oracle inequalities in risk minimization," *The Annals of Statistics*, vol. 34, no. 6, pp. 2593–2656, 2006.

[24] V. Koltchinskii, *Oracle Inequalities in Empirical Risk Minimization and Sparse Recovery Problems*. Springer, 2011.

[25] A. Agarwal, S. Negahban, and M. J. Wainwright, "Stochastic optimization and sparse statistical recovery: Optimal algorithms for high dimensions," in *NIPS*, pp. 1547–1555, 2012.

Appendix A. Proof of Lemma 3

From the exp-concavity of expected loss function we have

$$\mathcal{L}(\mathbf{w}_*) \geq \mathcal{L}(\mathbf{w}_i) + (\mathbf{w}_* - \mathbf{w}_i)^\top \nabla \mathcal{L}(\mathbf{w}_i) + \frac{\beta}{2}(\mathbf{w}_i - \mathbf{w}_*)^\top \mathrm{E}\left[[\ell'(y\mathbf{w}_i^\top \mathbf{x})]^2 \mathbf{x}\mathbf{x}^\top\right](\mathbf{w}_i - \mathbf{w}_*).$$

Combining the above inequality with our assumption

$$\mathrm{E}\left[[\ell'(y\mathbf{w}_i^\top \mathbf{x})]^2 \mathbf{x}\mathbf{x}^\top\right] \succeq \theta \mathrm{E}[\mathbf{x}\mathbf{x}^\top],$$

and rearranging the terms results in the following inequality

$$\mathcal{L}(\mathbf{w}_i) - \mathcal{L}(\mathbf{w}_*) + \frac{\theta\beta}{2}\|\mathbf{w}_i - \mathbf{w}_*\|_{\mathbf{H}}^2 \leq (\mathbf{w}_i - \mathbf{w}_*)^\top \nabla \mathcal{L}(\mathbf{w}_i).$$

Applying the fact that

$$\begin{aligned} & (\mathbf{w}_i - \mathbf{w}_*)^\top \nabla \mathcal{L}(\mathbf{w}_i) \\ = \; & \ell'(y\mathbf{w}_i^\top \mathbf{x})(\mathbf{w}_i - \mathbf{w}_*)^\top \mathbf{x} + (\mathbf{w}_i - \mathbf{w}_*)\left(\nabla \mathcal{L}(\mathbf{w}_i) - \ell'(y\mathbf{w}_i^\top \mathbf{x})\mathbf{x}\right) \\ = \; & \frac{\|\mathbf{w}_i - \mathbf{w}_*\|_{\mathbf{Z}_i}^2}{2\eta_i} - \frac{\|\mathbf{w}_{i+1} - \mathbf{w}_*\|_{\mathbf{Z}_i^{-1}}^2}{2\eta_i} + \frac{\eta G^2}{2}\mathbf{x}_i^\top \mathbf{Z}_i \mathbf{x}_i + (\mathbf{w}_i - \mathbf{w}_*)^\top \delta_i \end{aligned}$$

we obtain

$$\begin{aligned} & \mathcal{L}(\mathbf{w}_i) - \mathcal{L}(\mathbf{w}_*) + \frac{\theta\beta}{2}\|\mathbf{w}_i - \mathbf{w}_*\|_{\mathbf{H}}^2 \\ \leq \; & \frac{\|\mathbf{w}_i - \mathbf{w}_*\|_{\mathbf{Z}_i}^2}{2\eta_i} - \frac{\|\mathbf{w}_{i+1} - \mathbf{w}_*\|_{\mathbf{Z}_i}^2}{2\eta_i} + \frac{\eta G^2}{2}\mathbf{x}_i^\top \mathbf{Z}_i \mathbf{x}_i + (\mathbf{w}_i - \mathbf{w}_*)^\top \delta_i. \end{aligned}$$

Using the fact that $\eta_1 \geq 3/[\theta\beta]$,

$$\frac{\|\mathbf{w}_i - \mathbf{w}_*\|^2_{\mathbf{Z}_i}}{2\eta_i} = \frac{\|\mathbf{w}_i - \mathbf{w}_*\|^2_{\mathbf{M}_i}}{\eta_1}$$

and

$$\frac{\|\mathbf{w}_{i+1} - \mathbf{w}_*\|^2_{\mathbf{Z}_i}}{2\eta_i} = \frac{\|\mathbf{w}_{i+1} - \mathbf{w}_*\|^2_{\mathbf{M}_i}}{\eta_1} = \frac{\|\mathbf{w}_{i+1} - \mathbf{w}_*\|^2_{\mathbf{M}_i}}{\eta_1} - \frac{1}{2\eta_1}\left[(\mathbf{w}_i - \mathbf{w}_*)^\top \mathbf{x}_i\right]^2,$$

we get

$$\begin{aligned}
&\mathcal{L}(\mathbf{w}_i) - \mathcal{L}(\mathbf{w}_*) + \frac{\theta\beta}{3}\|\mathbf{w}_i - \mathbf{w}_*\|^2_{\mathbf{H}} \\
&\leq \frac{\|\mathbf{w}_i - \mathbf{w}_*\|^2_{\mathbf{M}_i}}{2\eta_1} - \frac{\|\mathbf{w}_{i+1} - \mathbf{w}_*\|^2_{\mathbf{M}_{i+1}}}{2\eta_1} + \frac{\eta_1 G^2}{2}\mathbf{x}_i^\top \mathbf{M}_i^{-1}\mathbf{x}_i + (\mathbf{w}_i - \mathbf{w}_*)^\top \delta_i \\
&\quad + \frac{\theta\beta}{3}\left(\left[(\mathbf{w}_i - \mathbf{w}_*)^\top \mathbf{x}_i\right]^2 - \|\mathbf{w}_i - \mathbf{w}_*\|^2_{\mathbf{H}}\right).
\end{aligned}$$

Appendix B. Proof of Lemma 4

Since $\mathbf{M}_i = \mathbf{M}_{i-1} + \mathbf{x}_i\mathbf{x}_i^\top$, we have

$$\mathbf{x}_i^\top \mathbf{M}_i^{-1}\mathbf{x}_i^\top = \mathrm{trace}(\mathbf{M}_i^{-1}(\mathbf{M}_i - \mathbf{M}_{i-1}) = \mathrm{trace}(\mathbf{I} - \mathbf{M}_i^{-1}\mathbf{M}_{i-1}).$$

Let $\gamma_j \geq 0, j = 1, \ldots, \ldots, d$ be the eigenvalues of $\mathbf{M}_i^{-1/2}\mathbf{M}_{i-1}\mathbf{M}_i^{-1/2}$. We have

$$\begin{aligned}
\mathrm{trace}\left(\mathbf{I} - \mathbf{M}_i^{-1}\mathbf{M}_{i-1}\right) &= \mathrm{trace}\left(\mathbf{I} - \mathbf{M}_i^{-1/2}\mathbf{M}_{i-1}\mathbf{M}_i^{-1/2}\right) \\
&= \sum_{j=1}^{d}(1 - \gamma_j) \\
&\leq \sum_{j=1}^{d}\ln\frac{1}{\gamma_j} \\
&= -\ln\det\left(\mathbf{M}_i^{-1/2}\mathbf{M}_{i-1}\mathbf{M}_i^{-1/2}\right) = \ln\det(\mathbf{M}_i) - \ln\det(\mathbf{M}_{i-1}),
\end{aligned}$$

which concludes the proof.

Appendix C. Proof of Lema 5

Define $\mathbf{V} = \mathbf{M}_n - \mathbf{M}_0 = \sum_{i=1}^n \mathbf{x}_i\mathbf{x}_i^\top$. Let $\lambda_j, j = 1, \ldots, d$ be the eigenvalues of $\mathbf{V}$. It is easy to verify that

$$\log\det(\mathbf{M}_n) - \log\det(\mathbf{M}_0) = \sum_{j=1}^{d}\log\left(1 + \frac{\lambda_j}{a}\right)$$

Since $\sum_{j=1}^d \lambda_j = n$, we have

$$\log\det(\mathbf{M}_n) - \log\det(\mathbf{M}_0) \leq \max_{\lambda_j \geq 0, \sum_{j=1}^d \lambda_j \leq n} \sum_{j=1}^{d}\log\left(1 + \frac{\lambda_j}{a}\right).$$

It is easy to verify that the above optimization takes its optimal at $\lambda_j = n/d, j = 1, \ldots, d$.

Appendix D. Proof of Lemma 6

The proof is based on the Bernstein inequality for martingales (see e.g., [21]).

Theorem 6. *(Bernstein's inequality for martingales). Let $X_1, \ldots, X_n$ be a bounded martingale difference sequence with respect to the filtration $\mathcal{F} = (\mathcal{F}_i)_{1 \le i \le n}$ and with $\|X_i\| \le K$. Let*

$$S_i = \sum_{j=1}^{i} X_j$$

be the associated martingale. Denote the sum of the conditional variances by

$$\Sigma_n^2 = \sum_{t=1}^{n} \mathrm{E}\left[X_t^2 | \mathcal{F}_{t-1}\right]$$

Then for all constants t, $\nu > 0$,

$$\Pr\left[\max_{i=1,\ldots,n} S_i > t \text{ and } \Sigma_n^2 \le \nu\right] \le \exp\left(-\frac{t^2}{2(\nu + Kt/3)}\right)$$

and therefore,

$$\Pr\left[\max_{i=1,\ldots,n} S_i > \sqrt{2\nu t} + \frac{\sqrt{2}}{3}Kt \text{ and } \Sigma_n^2 \le \nu\right] \le e^{-t}$$

Define martingale difference

$$X_i = \langle \mathbf{w}_i - \mathbf{w}_*, \nabla \mathcal{L}(\mathbf{w}_i) - \ell'\left(y_i \mathbf{x}_i^\top \mathbf{w}_i\right) \mathbf{x}_i \rangle$$

and martingale $\Lambda = \sum_{i=1}^n X_i$. Define the conditional variance Σ_n^2 as

$$\Sigma_n^2 = \sum_{i=1}^{n} \mathrm{E}_i\left[X_i^2\right] \le G^2 \sum_{i=1}^{n} \|\mathbf{w}_i - \mathbf{w}_*\|_{\mathbf{H}}^2 = G^2 A$$

Define

$$K = \max_i |X_i| \le 2RG.$$

Since $A \le 4R^2 n$, we have

$$\begin{aligned}
&\Pr\left(\Lambda \ge 2G\sqrt{At} + \sqrt{2}Kt/3\right) \\
&= \Pr\left(\Lambda \ge 2G\sqrt{At} + \sqrt{2}Kt/3, A \le 4R^2 n\right) \\
&= \Pr\left(\Lambda \ge 2G\sqrt{At} + \sqrt{2}Kt/3, \Sigma_n^2 \le G^2 A, A \le 4R^2 n\right) \\
&\le \Pr\left(\Lambda \ge 2G\sqrt{At} + \sqrt{2}Kt/3, \Sigma_n^2 \le G^2 A, A \le \frac{4R^2}{n}\right) \\
&\quad + \sum_{i=1}^{m} \Pr\left(\Lambda \ge 2G\sqrt{At} + \sqrt{2}Kt/3, \Sigma_n^2 \le G^2 A, \frac{2^{i+1}R^2}{n} < A \le \frac{2^{i+2}R^2}{n}\right) \\
&\le \Pr\left(A \le \frac{4R^2}{n}\right) + \sum_{i=1}^{m} \Pr\left(\Lambda \ge 4GR\sqrt{\frac{2^i}{n}t} + \sqrt{2}Kt/3, \Sigma_n^2 \le \frac{4R^2G^2}{n}2^i\right) \\
&\le \Pr\left(A \le \frac{1}{n}\right) + me^{-t}
\end{aligned}$$

where $m = \lceil 2\log_2 n \rceil$. The last step follows the Bernstein inequality for martingales. We complete the proof by setting $t = \ln(m/\delta)$ and using the fact

$$2G\sqrt{A} = \frac{6G^2}{\theta\beta} + \frac{\theta\beta}{6}A.$$

Appendix E. Proof of Lemma 8

$$\sum_{i=1}^{n}\mathcal{L}(\mathbf{w}_i)-\mathcal{L}(\mathbf{w}_*)\leq\sum_{i=1}^{n}\mathrm{E}_{(\mathbf{x},y)}\left[\langle\mathbf{w}_i-\mathbf{w}_*,\ell'(y,\mathbf{x}^\top\mathbf{w}_i)\mathbf{x}\rangle\right]-\frac{\theta}{2}\sum_{i=1}^{n}\langle\mathbf{w}_i-\mathbf{w}_*,\mathbf{H}(\mathbf{w}_i-\mathbf{w}_*)\rangle$$

Since

$$\begin{aligned}
&\sum_{i=1}^{n}\mathrm{E}_{(\mathbf{x},y)}\left[\langle\mathbf{w}_i-\mathbf{w}_*,\ell'(y,\mathbf{x}^\top\mathbf{w}_i)\mathbf{x}\rangle\right]\\
&\quad\leq\quad\sqrt{n}\sqrt{\sum_{i=1}^{n}\left(\mathrm{E}_{(\mathbf{x},y)}\left[\langle\mathbf{w}_i-\mathbf{w}_*,\ell'(y,\mathbf{x}^\top\mathbf{w}_i)\mathbf{x}\rangle\right]\right)^2}\\
&\quad\leq\quad G\sqrt{n}\sqrt{\sum_{i=1}^{n}\langle\mathbf{w}_i-\mathbf{w}_*,\mathbf{H}(\mathbf{w}_i-\mathbf{w}_*)\rangle}\\
&\quad=\quad G\sqrt{nA}\leq 2RG,
\end{aligned}$$

we obtain the desired inequality as

$$\sum_{i=1}^{n}\mathcal{L}(\mathbf{w}_i)-\mathcal{L}(\mathbf{w}_*)+\frac{\theta}{2}\sum_{i=1}^{n}\langle\mathbf{w}_i-\mathbf{w}_*,\mathbf{H}(\mathbf{w}_i-\mathbf{w}_*)\rangle\leq 2RG.$$